

Dimension Reduction of Hyperspectral Data on Reconfigurable Computers

Jacqueline Le Moigne, Pen-Shu Yeh,
Joanna Joiner
NASA Goddard Space Flight Center
Codes 935, 564, 910
Greenbelt, MD 20771

Tarek El-Ghazawi, Abhishek Agarwal
Dept Electrical & Computer Engineering
The George Washington University
801 22nd Street, NW
Washington, DC 20052

Gregory Donohoe
Electrical & Computer Engineering Dept
University of Idaho
Moscow, Idaho 83844-1023

Wei Xia
Global Science & Technology, Inc.
7855 Walker Drive
Greenbelt, MD 20770

Abstract- The objective of this work is to demonstrate the use of reconfigurable computing for on-board automatic processing of remote sensing data. The Field Programmable Processor Array (FPPA), a radiation tolerant alternative to Field Programmable Gate Arrays, developed at NASA/Goddard under ESTO funding, is the computation engine of our study, while preliminary feasibility studies are also performed on an hybrid reconfigurable architecture, the SRC-6. For this feasibility study, one of the first basic methodologies that is considered deals with dimension reduction of hyperspectral data. This processing represents one of the few most important first steps to be performed on-board, since it enables to reduce communication bandwidth, and make subsequent computations simpler, faster and more accurate. An adapted software environment developed for the FPPA is utilized for the design of this algorithm.

I. INTRODUCTION

The objective of this work is to demonstrate the use of reconfigurable computing for on-board automatic processing of remote sensing data. Enabling on-board data processing would reduce the data downlink bandwidth requirements by both preprocessing data and selecting data to be transmitted based upon predetermined content-based criteria. On-board processing would also reduce the cost and the complexity of a ground processing system so that it would be affordable to a larger community, including educational institutions. Other applications that would benefit from on-board processing are future reconfigurable web sensors missions as well as future Mars and planetary exploration missions, for which on-board processing would enable autonomous decisions to be taken on-board.

Among the computers capable of performing sophisticated computations, we focus on low-power and radiation-tolerant architectures. Among those, Field Programmable Gate Arrays (FPGAs) represent a flexible architecture that could be targeted for our purpose, but also presents the following weaknesses [1]:

1. the very fine granularity of FPGAs is not efficient, with only 1% of the chip available for logic and 99% used for interconnect and configuration. This usually results in a penalty in terms of speed and power.

2. high-performance FPGAs are difficult to radiation-harden, and currently-available rad-tolerant FPGAs have two orders of magnitude fewer equivalent gates than commercial FPGAs.

In the light of these considerations, we are considering a radiation tolerant alternative to FPGAs that was developed at NASA Goddard under a previous award, the Field Programmable Processor Array (FPPA): this architecture will be described in section III.A. As a reference when performing experimental feasibility studies, we are also using an hybrid reconfigurable architecture, the SRC-6E, described in section III.B. The SRC-6E could not be flight-qualified as is, but it can be utilized as proof-of-concept, prior to FPPA implementation.

To test the feasibility of using reconfigurable computers to perform onboard processing, we are targeting a few applications, and the first one described in this paper deals with hyperspectral data reduction. With hundreds or thousands of channels, hyperspectral imagery possesses much richer spectral information than multispectral imagery. In particular, hyperspectral imaging has a wide range of applications in mining, geology, forestry, agriculture, and environmental management. However, realizing the full potential of hyperspectral technology remains a challenge. It is clear that more effective data processing techniques are needed to deal with hyperspectral cubes. One example is land image classification with sensors such as AVIRIS (224 bands). For this application, it is necessary to have a minimum ratio of training pixels to the number of spectral bands (in order to ensure a reliable estimate of class statistics), and as a consequence, dimension reduction has become a significant part of hyperspectral land classification. Another example is the Atmospheric Infrared Sounder (AIRS) instrument suite which is designed to measure the Earth's atmospheric water vapor and temperature profiles on a global scale. With 2378 channels, the AIRS Infrared data represent a good candidate for dimension reduction, as described in section II.B.

In this paper, we will show how a newly-developed dimension reduction based on wavelets can be implemented on reconfigurable computing architectures such as the FPPA and the SRC-6E. We will also show how this method, already tested with sequential algorithms on AVIRIS data, can be extended to perform wavelet analysis reduction method of Atmospheric Infrared Sounder (AIRS) data. Preliminary results will be presented in section IV.

II. DIMENSION REDUCTION OF HYPERSPECTRAL DATA

A. Wavelet-Based Dimension Reduction

As a definition, data or dimension reduction is a process designed to reduce data volumes by filtering out specific redundant information. It is also commonly used for performing image fusion. Such a feature vector dimensionality reduction has been pursued in several different ways, but traditionally, Principal Component Analysis (PCA) has been the technique of choice.

Recently, we developed a new wavelet-based data reduction method [2] and implemented it both sequentially and on a Commercial-Off-The-Shelf (COTS)-based architecture, the Beowulf. The principle of this novel wavelet-based method is to apply a discrete one-dimensional wavelet transform in the spectral domain and at each pixel. This transform decomposes the signature of each pixel into a set of composite bands that are linear, weighted combinations of the original spectral bands. Figure 1 shows an example of the actual signature of one class (Corn) for 192 bands of an AVIRIS hyperspectral dataset, and different levels of wavelet decomposition of this

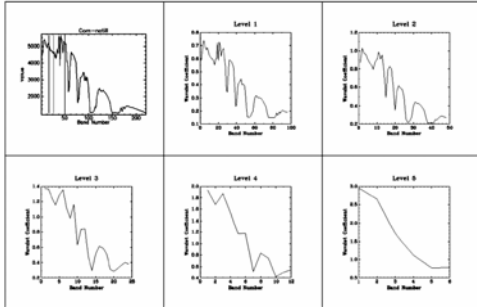


Figure 1 - Example of the Corn Spectral Signature And Different Levels of Wavelet Decomposition

spectral signature. When the number of bands is reduced, the structure of the spectral signature becomes smoother than the structure of the original signature, but the signal still shows the most important features for several levels. The detailed description of the algorithm is then:

- (1) Multi-resolution wavelet decomposition of each pixel 1-D spectral signature.
- (2) At each level of decomposition:
 - (2.1) Reconstruction using only low-pass information.

- (2.2) Similarity measure (e.g., correlation) between original signature and reconstructed signature for that decomposition level.
 - (2.3) Record that level in a histogram if it satisfies a quality of "good" reconstruction, defined by a percentage-threshold.
- (3) From the histogram, choose the optimum level of decomposition. and build the corresponding dimension reduced image.

From a complexity point of a view, the whole algorithm complexity is in the order of $O(MN)$, where M is the number of pixels in the spatial domain and N is the number of bands. The total estimated complexity of PCA is $O(MN^2+N^3)$, which shows that the computational efficiency of the wavelet reduction technique is superior to the efficiency of the PCA method.

We validated the wavelet-based reduction using different supervised classifications of hyperspectral AVIRIS data. The results show that our new wavelet-based dimension reduction method provides a greater computational efficiency as well as a better or comparable overall classification than the widely used PCA method [2]. In this paper, we will introduce the implementation of the wavelet-based dimension reduction algorithm on the FPPA and on the SRC-6E, as well as the generalization of this algorithm for reducing the hyperspectral AIRS sounder data.

B. Application to AIRS Data

The Atmospheric Infrared Sounder (AIRS) is a high-spectral-resolution IR spectrometer designed to answer some of NASA and NOAA research objectives in weather prediction [3]. AIRS provides more accurate global temperature and moisture profiles than previous instruments, which are required to improve short-range weather forecasts. Due to array technology in HgCdTe detectors, it is now possible to obtain contiguous coverage of most of the 3.7-15.4 μm region of the spectrum with more than 2000 spectral samples. This hundred-fold increase in the number of channels of AIRS compared to previous similar instruments is becoming a computational challenge for ground data system [3]. Generally, AIRS data is reduced depending on the application by choosing "appropriately" the channels to process. In order to optimize the information content kept in the processing, better and automatic data reduction algorithms are needed that would reduce the amount of data to process while keeping data accurate and meaningful.

III. RECONFIGURABLE HARDWARE

A. Field Programmable Processor Array (FPPA)

A new reconfigurable processor called the Field Programmable Processor Array (FPPA) is being applied to the wavelet-based dimension reduction algorithm. The FPPA is a derivative of the Reconfigurable Data Path

Processor (RDPP) program funded by ESTO. Like the RDPP, the FPPA is a reconfigurable processor chip with 16 on-board processing elements, programmable interconnect, and an on-board execution unit. It implements a synchronous data flow computational model, particularly well-suited to processing streaming data from spacecraft instruments. The FPPA differs from its predecessor in a few details. It employs 16 bit data paths, instead of 24 bits; this enables it to operate faster and with lower power consumption. Unlike the RDPP, the FPPA supports multiple-precision arithmetic. This removes any limitations imposed by the 16-bit data path, and makes the chip adaptable to a wider range of applications.

(1) Reconfigurable Platform

The FPPA chip will serve as a computational accelerator, but it must operate in the context of a system. The FPPA team is developing a computational platform, which will incorporate 16 FPPA chips (with a total of 256 processing elements), configurable memory blocks, configurable interconnect, and an on-board microcontroller to serve as host and to coordinate the operations of the other components.

(2) FPPA Design Software

The FPPA team has developed a suite of software tools to program the chip. These include a simulator, floating point to fixed point conversion tools, a configuration assembler, and a run-time assembler. The assemblers require the programmer to write configuration and run-time

specifications in a processor-specific language. The dataflow architecture, which is unfamiliar to most programmers, and the special-purpose language, present a challenge to someone learning to program the chip. In particular, the dataflow process, which is easy to comprehend graphically, is difficult to capture in code. To facilitate programming the FPPA, the design team has created a graphical front end based on Simulink. The programmer selects processing blocks from an FPPA menu, places them on a work screen, and connects inputs to outputs with by drawing lines. Each block can be “opened” to allow text entry of parameters.

Figure 2 shows the programmer’s view of the interface, implementing a 4-tap finite impulse response (FIR) filter. This is the view inside one FPPA chip. The colored rectangles represent processing elements. The first is configured to perform a multiply operation, and the last three, to multiply and add. The squares with numbers inside represent constants, stored in constant registers within the processing elements. The design can be simulated by clicking the button to execute a Simulink simulation. Currently, the interface invokes a Matlab simulation that is architecturally equivalent to the FPPA or FPPA platform. The next step is to automatically compile FPPASim code and to invoke the FPPA simulator for a bit-accurate simulation.

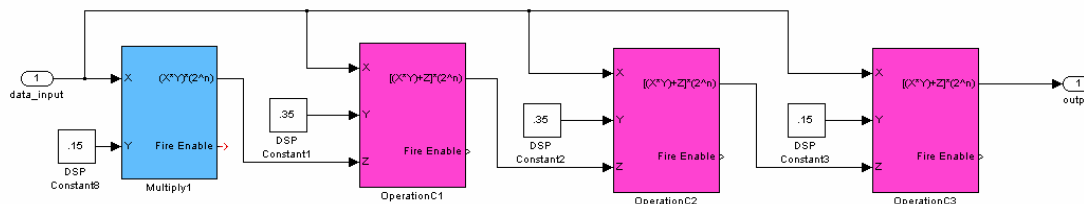


Figure 2
Simulink interface to FPPA software

B. Hybrid Reconfigurable System, SRC-6E

(1) Hardware Architecture

The SRC-6E platform consists of two general-purpose microprocessor boards and one MAP® reconfigurable processor board. Each microprocessor board is based on two 1 GHz Pentium 3 microprocessors. The SRC MAP board consists of two MAP reconfigurable processors. Overall, the SRC-6E system provides a 1:1 microprocessor to FPGA ratio.

Microprocessor boards are connected to the MAP board through the SNAP interconnect. SNAP card plugs into the DIMM slot on the microprocessor motherboard. Hardware architecture of the SRC MAP processor is shown in Fig. 3.

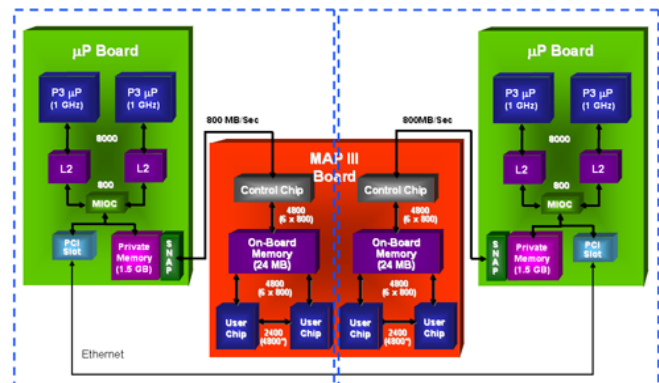


Figure 3
Hardware Architecture of SRC-6E

This processor consists of two programmable User FPGAs, six 4 MB banks of the on-board memory (OBM), and a single Control FPGA. In the typical mode of operation, input data is first transferred through the Control FPGA from the microprocessor memory to OBM. This transfer is followed by computations performed by the User FPGA, which fetches input data from OBM and transfers results back to OBM. Finally, the results are transmitted back from OBM to microprocessor memory.

(2) Programming Model

The SRC-6E has a similar compilation process as a conventional microprocessor based computing system, but needs to support additional tasks in order to produce logic for the MAP reconfigurable processor, as shown in Fig. 4.

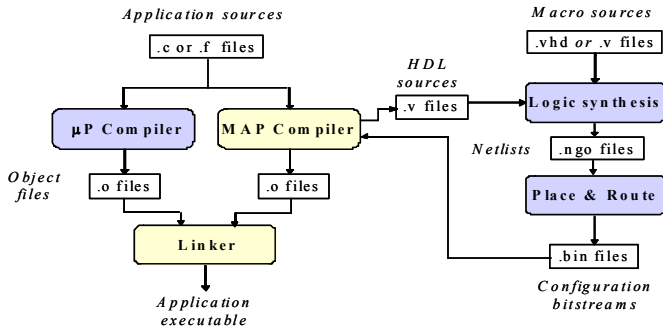


Figure 4
SRC Compilation Process

There are two types of application source files to be compiled. Source files of the first type are compiled targeting execution on the Intel platform. Source files of the second type are compiled targeting execution on the MAP reconfigurable processor. A file that contains a program to be executed on the Intel processor is compiled using the microprocessor compiler. All files containing functions that call hardware macros and thus execute on the MAP are compiled by the MAP compiler. MAP source files contain MAP functions composed of macro calls. Here, macro is defined as a piece of hardware logic designed to implement a certain function. Since users often wish to extend the built-in set of operators, the compiler allows users to integrate their own VHDL/Verilog macros.

IV. PRELIMINARY RESULTS

A. Wavelet Decomposition on the FPPA

The wavelet decomposition consists of convolution with a series of fixed kernels (that is, a finite impulse response or FIR filter), followed by decimation or subsampling. A chain of these operations produces a set of signals from which the original can be reconstructed.

For this program, the 4-weight Daubechies “DAUB4” wavelet basis set was used, using only the low-pass filter components to reduce the dimensionality of hyperspectral signal vectors. Each level uses only four filter weights, which enables us to implement four decomposition filters in

one FPPA chip. The average throughput is one data sample per instruction cycle, with a 4-sample latency between successive samples, as the output of one filter serves as the input to the next.

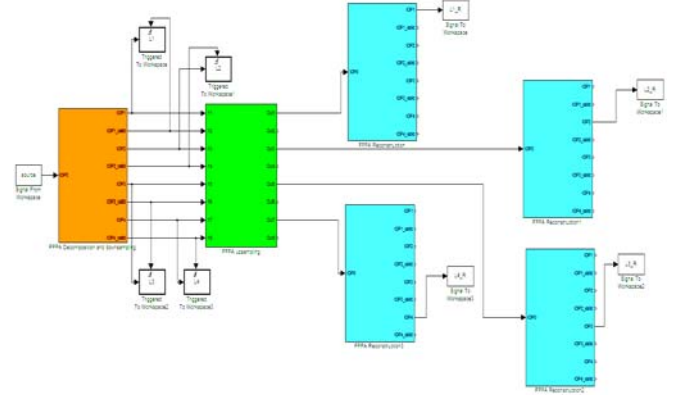


Figure 5
Programmer's View of Wavelet Decomposition and Reconstruction Using Five FPPA Chips

To demonstrate the correctness of the wavelet decomposition, the FPPA simulator was used to reconstruct the original signals from different decomposition levels. These two processes – decomposition and reconstruction – represent two key operations for wavelet-based dimensionality reduction. Figure 5 shows the programmer's view of the wavelet decomposition and reconstruction.

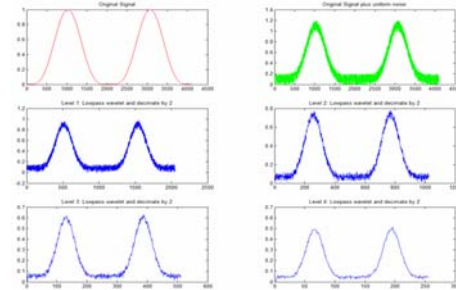


Figure 6
Gaussian Signal with Additive Noise and 4 Levels of Wavelet Decomposition

This implementation requires five FPPA chips, represented by the large rectangles. The chip on the left computes the DAUB4 low-pass wavelet decomposition. The other four chips perform up-sampling and reconstruction. Each chip reconstructs the signal from one of the decomposition levels. Figure 6 shows a test signal made up of a pair of Gaussian signals (top left) with additive white noise (top right). The next four images are four different DAUB4 wavelet decompositions.

B. Wavelet-Based Dimension Reduction on the SRC-6E

An initial fully pipelined implementation of the automatic wavelet dimension reduction algorithm has been developed on the SRC-6E. The design is complete in the functional sense, although many improvements are still in the plan.

The improvements will enhance the modularity of the architecture, and might improve the speed of processing as well.

The current implementation has shown an order of magnitude improvement of speed over a 1.8GHz Xeon P4 based PC. The implementation of this kernel allowed us to exploit the different features and limitations of reconfigurable computers by experimenting with a state-of-the-art reconfigurable computer, namely SRC-6E. Two main implementations are developed to understand all design issues and tradeoffs and assess the hardware design approach benefits. These are pure C running on a P4 and P3 microprocessor, and pure hardware design in VHDL with an overlap and without overlap of I/O with processing. The experiments were run over an AVIRIS called Salinas'98. This AVIRIS dataset was acquired on October 9, 1998, South of the city of Greenfield in the Salinas Valley in California. The dataset was taken at low altitude with a pixel size of 3.7 m. It consists of 217x512 pixels by 192 bands of radiance data.

From these experiments we have learned a great deal about the opportunities to maximize concurrency in the underlying algorithm. We have therefore developed a formal technique for optimizing the performance of a reconfigurable computer. A mathematical model for this approach has been derived, for a generic reconfigurable machine, taking into account the parameters and constraints imposed by both the system and the application. This technique depends on overlapping the computations on the User FPGAs with the I/O transfer. This overlapping requires dividing data transfers into multiple partial transfers that can be overlapped with partial computations.

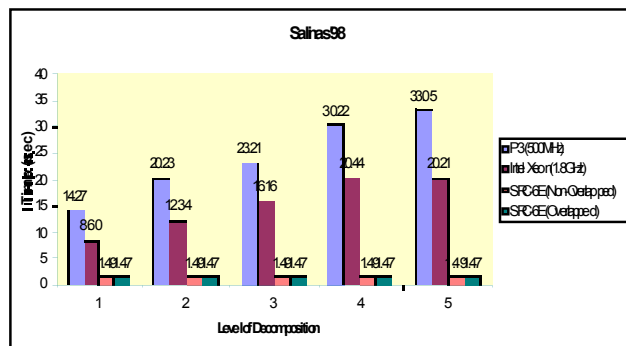


Figure 7 - Execution Time

C. Wavelet-Based Dimension Reduction of AIRS Data

The dataset used for this study was acquired on January 1, 2003. The AIRS Infrared (IR) level 1B data set contains AIRS infrared calibrated and geolocated radiances in milliWatts/m**2/cm**-1/steradian. This data set is generated from AIRS level 1A digital numbers (DN), which includes 2378 infrared channels in the 3.74 to 15.4 μ m region of the spectrum. A day's worth of AIRS data is

divided into 240 scenes each of 6 minutes duration. For the AIRS infrared measurements, an individual scene consists of 135 scan lines containing 90 cross-track footprints; thus there is a total of 135 x 90 = 12,150 footprints per AIRS IR scene. The 2378 infrared channels have been broken into three major intervals:

- 3.74 - 4.61 micron, channel 1865 – 2378
- 6.20 - 8.22 micron, channel 1263 – 1864
- 8.80 - 15.4 micron, channel 1 – 1262

Each of these intervals is further divided into subintervals for a total of 17 subintervals, not equally spaced.

In order to adapt to this configuration of the data, the automatic wavelet-based dimension reduction code was modified and applied both considering 3 or 17 intervals.

3 intervals ; Original # of Channels = 2378				
Interval	Threshold	Optimum Level of Decomposition from the Histogram	Corresponding Reduction	Reduced # of Channels
1	98%	1	50%	1189
2	95%	1	50%	1189
3	92%	1	50%	1189

The results are the following:

a) Using three major intervals, the results show that the 3 intervals do not produce similar reduction at various correlation thresholds. To get the same level of reduction, thresholds have to be chosen differently as shown above.

b) Using 17 sub intervals, we observed the following:

17 Sub-Intervals		Threshold = 98%		Threshold = 92%	
Sub-Interval	Original # Channels	Optimum Level Chosen	Reduced # Channels	Optimum Level Chosen	Reduced # Channels
0	118	0	118	1	59
1	116	0	116	2	29
2	130	1	65	3	16
3	150	0	150	0	150
4	104	1	52	3	13
5	106	1	53	3	13
6	192	0	192	0	192
7	94	0	94	2	23
8	106	0	106	2	26
9	159	0	159	0	159
10	167	0	167	0	167
11	167	0	167	0	167
12	161	0	161	0	161
13	167	0	167	0	167
14	167	0	167	0	167
15	144	2	36	4	9
16	130	0	130	0	130

These results show that different sub-intervals have different properties, including anomalies, that yield different optimum levels for dimension reduction. Currently, scientists choose channels by looking at low

correlation for certain intervals and high correlation for other intervals. This selection of bands depends on prior information about the information contained in the interval. We are currently modifying the code to include this information.

V. CONCLUSION AND FUTURE WORK

To test the feasibility of using reconfigurable computers to perform onboard processing, we performed experiments on 2 different reconfigurable architectures for the purpose of hyperspectral data reduction. We demonstrated that wavelet decomposition and reconstruction can be mapped on the FPPA architecture. Future work will study the implementation of the full wavelet-based dimension reduction on the FPPA. We also prototyped the automatic wavelet-based dimension reduction algorithm over the SRC-6E architecture. We observed a 10x speedup using the P3 version of SRC-6E. From our previous experience we expect this speedup to double using the P4 version of SRC machine. These speedup figures were obtained while I/O is still dominating. This speedup can be increased by using faster I/O techniques.

Furthermore, we generalized the wavelet-based dimension reduction algorithm for the 2378 bands-AIRS data. Results showed that additional information content about the data needs to be incorporated, and this will be studied first at the level of the correlation, second in the use of the high-pass information provided by the wavelet decomposition.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the support of the Earth Science Technology Office (ESTO) Advanced Information Systems Technology Program.

VI. REFERENCES

- [1] A. DeHon, "Reconfigurable Architectures for General-Purpose Computing," MIT AI Tech. Rep. No. 1586, October 1996.
- [2] S. Kaewpijit, J. Le Moigne, and T. El-Ghazawi, "Automatic Reduction of Hyperspectral Imagery Using Wavelet Spectral Analysis," IEEE Transactions on Geoscience and Remote Sensing, Vol.41, No.4, pp.863-871, April 2003.
- [3] H. Aumann, R. Pagano, "Atmospheric Infrared Sounder on the Earth Observing System," *Optical Eng.*, 33, No. 3, 1994.